

Email overload: exploring personal information management of email

Steve Whittaker

Lotus Development Corporation

One Rogers St

Cambridge MA. 02142, USA

+1 (617) 693 5003

Steve_Whittaker@crd.lotus.com

Candace Sidner

+1 (617) 693 7737

Candy_Sidner@crd.lotus.com

ABSTRACT

Email is one of the most successful computer applications yet devised. Our empirical data show however, that although email was originally designed as a *communications* application, it is now being used for additional functions, that it was not designed for, such as *task management* and *personal archiving*. We call this *email overload*. We demonstrate that *email overload* creates problems for personal information management: users often have cluttered inboxes containing hundreds of messages, including outstanding tasks, partially read documents and conversational threads. Furthermore, user attempts to rationalise their inboxes by filing are often unsuccessful, with the consequence that important messages get overlooked, or "lost" in archives. We explain how *email overloading* arises and propose technical solutions to the problem.

Keywords

Email, information overload, personal information management, asynchronous communication, filing, task management, interpersonal communication, ethnography, empirical studies.

WHY STUDY EMAIL?

Email is one of the most successful computer applications yet devised. There are millions of email users world-wide who often spend significant proportions of their work time using email. Research suggests that email has contributed to the growth of distributed organisations, by allowing people in different geographical areas to communicate across time and space. It has also led to the emergence of on-line communities by supporting asynchronous communication [8]. Email has been the subject of many studies, including pioneering early work that focussed on the social and communicative aspects of email, comparing its usage with face-to-face communication [8]. Nevertheless, there is little systematic data on its usage and utility as a workplace technology. Furthermore, the success and popularity of email has led to high daily volumes of email being sent and received. Research has not yet addressed how people organise and manage large amounts of information. This study therefore presents a quantitative analysis of

the mailboxes of 20 users, along with 34 hours of interviews to address these new questions.

Email applications were originally designed for *asynchronous communication*, but as our analysis will show, email has evolved to a point where it is now used for multiple purposes: document delivery and archiving; work task delegation; and task tracking. It is also used for storing personal names and addresses, for sending reminders, asking for assistance, scheduling appointments, and for handling technical support queries. We use the term *email overload* to describe the use of email for functions that it was not designed for. We discuss three main email functions: *task management*, *personal archiving* and *asynchronous communications*. The central question is how well a single tool can support all these functions. Subsidiary questions must also be asked in each category.

Task management requires users to ensure that information relating to current tasks is readily *available*. This both preserves *task context* and allows users to determine the progress of ongoing tasks. Task management also involves *reminding* oneself about when particular tasks or actions have to be executed [1,2,4,6]. How do people do this in email?

Personal archiving or *filing* addresses how people organise and categorise longer term information, so that it can later be retrieved. Archives are not of immediate relevance to current tasks, but are constructed for reference or anticipated future use. Research shows that users experience major problems in generating appropriate folder labels when *filing* longer term information for later retrieval, and in *reconstructing* these labels when they engage in later retrieval [1,2,4,6]. To what extent do these problems occur in email?

Asynchronous communication is concerned with interaction in a permanent medium across space and time. Research has characterised face-to-face workplace communications as consisting of repeated brief communications [3,10]. Such interactions are seldom one-shot, and workers often engage in multiple intermittent interactions in order to complete a task. Workers are also usually engaged in several independent, but concurrent ongoing conversations, with the requirements of tracking separate conversational threads and switching contexts between conversations [10]. Does email communication have these characteristics, and how are *asynchronous communications* conducted?

To provide preliminary answers to these questions, this study presents qualitative and quantitative information about the use of email for *task management*, *personal archiving* and *asynchronous communication*. We describe the problems people experience with

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

CHI 96 Vancouver, BC Canada
© 1996 ACM 0-89791-777-4/96/04..\$3.50

each of these functions, and the strategies they invoke to address the problems. Finally we suggest potential technical solutions.

SYSTEM AND METHOD

We studied users of NotesMail, the email component of Lotus Notes. This client-server system has a GUI with a number of standard features, including the ability to compose, reply to, copy (cc), and blind copy (bcc) messages to other users. Incoming unread messages are delivered into the inbox (called "uncategorised" in Notes), where they appear in a different colour from messages the user has "opened" and read. Once "opened", they appear in the standard colour. The system provides users with the ability to file information: they do so by creating categories (equivalent to folders in other email systems), so that related messages can be stored and accessed together, by category label. There is no conversational threading of messages: although responses generated using the reply option appear with *re* in the subject line, the system does not allow the user to automatically view messages from a given thread together. It is, however, possible for users to view messages in terms of various other properties. One can view by sender, date, size, or category. The system also has a text search capability: email databases can be full text indexed, so that key word or Boolean queries are possible. There is no formal support, or policy for, archiving.

The 20 study participants were office workers representing four major job types: 4 high level managers, i.e. people who had other managers reporting to them; 5 first level managers; 9 professional workers with no management responsibility, and 2 administrative assistants. All participants were experienced users with between 2 and 15 years experience of email. They had all used NotesMail for more than one year. All participants were employees of Lotus Development Corporation, a software development firm. We chose our subjects because we wanted to investigate email use in multiple job types, with different responsibilities. Our participants were drawn from marketing, consultancy, software development, support and research groups. We chose the organisation both because of its pervasive use of email and ready access to subjects. Given these choices, we studied NotesMail because it was the most frequently used system in the Lotus organisation.

We collected quantitative data about the mailbox of 18 users: (a) total number, age, and size of messages in their mailbox; (b) number of messages in each archival folder; (c) conversational threads. Due to technical difficulties, we were unable to collect quantitative data for the final two users. Ideally we would have wished to collect longitudinal data over an extended time period, to look at changes in mailbox size and structure, but the logistics of repeated access to personal data prevented this. We were however, able to collect a "snapshot" of each mailbox at a given point, from which we drew important inferences, which we report below.

We also interviewed all 20 participants for 1-2 hours using semi-structured questions. We asked them to describe: (a) the volume of email they sent and received; (b) their prioritisation, reading and reply strategies; (c) their correspondence management, e.g. when they used *reply*, *cc* and *bcc* features and how they managed conversational threads; (d) their filing behaviours. We also discussed: (e) the main problems they were experiencing with email; and (f) their reactions to certain technical solutions addressing these problems. Interviews were carried out in people's offices and

participants were encouraged to demonstrate their statements and strategies with reference to their actual systems.

We analysed our interviews by collecting user comments about each issue described above. We present representative quotes from participants about these, and where there was substantial disagreement or inconsistency between participants' opinions, then we present quotes representing alternative points of view.

EMAIL OVERLOAD: THE PROBLEM

Participants were generally highly positive about email as a communication tool. They stressed how it enabled them to collaborate with others across time and distance [8]. They also pointed out its advantages over other technologies such as the telephone, and even face-to-face interaction. Nevertheless certain individuals experienced major problems in reading and replying to email in a timely manner, with backlogs of unanswered email, and in finding information in email systems. The inability to effectively manage communication means lost information, and reduced responsiveness. These have clear negative outcomes for both individual and corporate productivity.

"Waiting to hear back from another ... employee can mean delays in accomplishing a particular task, which can ... have significant impact on our overall operations. Depending on the situation, it can either be critical or just frustrating."

"One of my pet-peevs is when someone does not get back to me, but I am one of the worst offenders. I get so many e-mails (average 30-40/day) and phone messages (15-20) that I cannot keep up and also do my real job..."

"given the sheer volume of stuff that passes through here. I mean I couldn't even give you a percentage of how much is missed. I mean -- not necessarily missed but certainly recorded but never followed up on"

"I dedicate somewhere between minimally two hours at the outlying range, up to ten hours on any given day trying to stay on top of email"

So why do these problems arise? A simple *one-touch* model of email might assume: incoming messages that are *informational*, i.e. those not requiring a response, are read, and then either deleted or filed, depending on their relevance. Incoming messages that form part of a *correspondence*, (i.e. requiring a response), are answered, and then either deleted or filed. According to the *one-touch* model, information can therefore be in two possible states: unread and filed. The user's inbox at any point should solely consist of a small number of unread incoming messages, and the rest of their mailbox consist of filed items.

Our quantitative data show the *one touch* model is patently incorrect. The mean number of inbox items is 2482, and the mean number of filed items (858) is small compared with the number of inbox items, so that the inbox constitutes on average 53% of people's mailfiles. It is implausible that users receive 2482 new items each day, so what is happening and why is the inbox so full? It turns out that there are two related reasons for this: (a) the inbox operates as a *task manager*, where people are *reminded* of current tasks, and where people can keep information relevant to those tasks *accessible*; (b) people find it hard to *file* information to remove it from their inbox, both because

filing into folders is difficult and there may also be few benefits to creating folders.

EXPLAINING THE FULL INBOX: MANAGING WORKING INFORMATION

Our users received a large number of messages each day (mean 49). One reason for the large volume of incoming messages is that email is now the source of many different office tasks, serving as the place in which work is received and delegated.

"it's where things come into your life in a way. It's the place where ... people hand things off to you, it's the electronic office"

Email can be an important determinant of how people spend their working day, again suggesting that it is a place that users receive and hand off tasks:

"I check it before I leave the house just in case there's anything I didn't get the night before. I read it as soon as I get into the office. It does change what I do throughout the day, like what – I may come in thinking I was going to do one thing, but get mail that sort of diverts me into doing something else.... If I haven't checked my mail it makes me uncomfortable. And there's invariably a piece of something I was supposed to do, that's time sensitive".

Both the volume of incoming mail and the fact that mail is being used for task management, leads to a breakdown in the *one-touch* model. While there is evidence that users try and process information at once, there are a number of reasons why immediate responses are sometimes not possible or not appropriate, so that incoming messages remain undischarged in the inbox.

One general reason relates to the amount of time users have currently available. If the message requires more than a certain amount of time or effort to process, then users delay dealing with it and proceed to other potentially more urgent or manageable messages in the inbox. There are also specific types of messages that are often not discharged immediately:

(a) *"To dos"* These are messages which require the user to execute some action. In some cases, the message may require the user to engage in further complex activities which might take days to achieve. The user does not usually suspend the process of reading email to discharge these activities. These *"to dos"* are commonly kept in the inbox as *reminders* of unfinished tasks.

(b) *"To reads"* Alternatively, messages might be long documents. Although these are often informational and do not require a reply, they still take time and effort to read, and users often delay reading them, so that the inbox may contain unread or partially read documents. The quantitative data support this. We found that on average 21% of the inbox (i.e. an average of 334 messages) were long, when a long message was defined as more than 10Kbytes (~5 screensfull).

(c) *Messages of indeterminate status* One issue with informational messages, is that users are often unsure of the significance of an incoming message when it first arrives. Rather than investing valuable time in reading it at once, they register its arrival, but delay dealing with it until some later point when they are more certain of its importance. What makes immediate decisions difficult, is that the value of a given message may depend on events that occur *after* the message has been received: a flurry of subsequent messages on the

topic may reveal its importance, or else it may turn out to be a "dead-end" with no follow-up being necessary. Rather than delete it immediately, users often conservatively retain it just in case it turns out to be important. This user describes keeping such documents in her inbox:

"I've gotten messages, I haven't dealt with them, haven't known what to do about them. And starting this new position, I got a lot of mail messages, but I didn't have the knowledge to know what on earth they were about ... You know, people were talking about servers and infrastructure...I just realized some day I would understand it, and I saved a lot of that stuff, and actually had to make a presentation where it all came in handy – where I went, I've been sitting with this information all this time and I didn't know it.... so you're conservative about keeping stuff, and ...at some point in the future it may come in useful."

People also explained how retaining such unread messages was useful for unplanned contingencies, e.g. when they received surprise phone calls about them. They were able to read the documents in the course of the call, while pretending to have already read them.

"[email] is the best that's ever happened for covering your ass ... because while the guy's on the phone, he says, well, I sent that to you a week and a half ago, and you think, shit, I never saw that. But you say, 'really? Yeah, oh yeah, I remember reading that', as you're reading it."

(d) *Ongoing correspondence*. Finally the inbox is sometimes used for ongoing, but incomplete, threads of asynchronous conversations. The user may delay responding to a question from another person because a careful reply is necessary which takes more time than is currently available. Alternatively, users may be unable to reply immediately, because they currently do not have the answer, and they await further information from other people.

For more complex interactions involving multiple exchanges over an extended time period, users may also track and sometimes save, both their own and other people's contributions to the conversation. An issue may take several email exchanges to be resolved, or users may require the responses of multiple individuals in order to collate opinion, or reach consensus. A major problem with email correspondence is that there are no agreed conventions about whether to include the context or history of prior messages as the conversation proceeds. Often this context is important, because it is necessary for interpreting what each subsequent message means. One user describes this problem with one of his coworkers.

"X is unbelievable in that he never puts the context in which, of what he's replying to. He always comes up with these one line responses, and I have no idea what it is that he's talking about, you know, it's like, "Re: the Internet." ... And, in many cases, he's replying to something that somebody else sent, but I wasn't on the original distribution list, but he thinks that I'd be interested ...So I get this totally out-of-context great idea"

For multiple interchanges extending over periods of days or weeks, it is easy to forget who said what and to whom. For legal purposes, it can also be important on occasion to record exactly what was said by whom. For these reasons, for certain critical interactions, users

sometimes save both the originating message, as well as subsequent interactions.

"the people that I consider some of the best problem-solvers in my organization are fanatical about the history. And you get the whole thing. Yeah, the audit trail of what happens. We do that, have to do that a lot in support... Because you could be dealing with an extended customer issue that bounced around... especially the people that are right on the front lines, it's almost like second nature to them... It's not only everything that's being said... It's every person that has been involved"

This conversational record serves multiple functions: an archive of what has been discussed; a *reminder* to the user that the conversation is in progress; and a record of the status of the conversation, and whether one "owes" or is "owed" a response. The importance of conversational tracking, as well as the absence of conventions about whether others will include history, can mean that each exchange of a lengthy conversation will appear as a separate message in the inbox. Not only does this increase the number of messages in the inbox, it is often difficult to gather together the related threads of a conversation, because conversational exchanges are often interleaved with other unrelated information.

"That reply with the history of previous stuff.... when some people do that and some don't and the fact that it's all interspersed with all other kinds of crap in my e-mail, and then I just can't pick up an e-mail and find out what else it belongs with".

The quantitative data indicate the pervasiveness of inbox conversational threads. We examined the subject line of each message and counted the number of instances in which it contained "re", signalling it was a reply to an original message. By this metric¹, we found that inboxes contain a mean of 209 such messages, and that these constitute 12% of the total inbox messages.

To summarise, multiple types of items linger in users' inboxes: actions the user has yet to do; documents that are partially read or unread, and correspondence that is still in progress. What unifies these is that they are all *incomplete*, and the usual strategy is to leave them in the inbox to serve as *reminders* that some further action is required. They are not normally filed away, because filing would mean that they are no longer visible whenever new email is read or the inbox searched.

"the reason that I don't categorize things but leave them in here [the inbox] is that I realize there are a certain number of things that I keep in my mind, and I will go back for... And others, I do have to count on tripping over them. And as long as there is that mess that I know I have to do the multiple passes of reading over... I'm kind of depending on that serendipitous tripping over it again as a way to remind me"

The importance of this visual reminding function is evidenced by the fact that five users had experimented with a strategy of filing undischarged information in an "action", or "to do", folder. In all but one of these cases, this action folder was abandoned, because users had to *explicitly remember* to go to it, open it and view its contents, rather than being reminded of these unintentionally, by the mere fact of being in the inbox reading new email².

"I used to have an "unread" folder. Which was messages I'd opened up, but I had never finished reading. Like those big ones that ... I didn't get through right away. I didn't go back to it [the unread folder] often enough though".

The single person who was successfully using "to do" folders had reconfigured her mailbox UI, so as to be reminded about this folder. Her "to do" folder appeared immediately above the first new unread item in her inbox, so that she would automatically see the "to do" folder whenever she read new email.

A second reason for leaving information in the inbox concerns its *availability*. In the case of extended interactions, users often keep conversational history in the inbox, because they believe it to be more *accessible* there.

"you may not want to file it. Because it might be something you need to refer to. I don't want to file that yet, because it's active ... there are things that are happening as a result of that. It's easier for me to find it. So I want to keep in my "in" box, keep it current"

GETTING INFORMATION OUT OF THE INBOX: THE PROBLEM OF FILING

We have seen that incomplete tasks being "kept around" can lead the inbox to be full. A second complementary reason is that users find it difficult to move messages out of the inbox by filing them into folders containing collections of related messages. Why is this? Firstly generating and maintaining these folder collections requires considerable effort. Secondly, and more importantly, the resulting collections may be of little use in message retrieval.

Filing is a cognitively difficult task [2,4]. Successful filing is highly dependent on being able to imagine future retrieval requirements. It is hard to decide which existing folder is appropriate, or, if a new folder is needed, how to give it a memorable name.

"any piece of information longer than five lines has at least several axes along which you might want to look it up and it really depends how you're coming at it and what you're thinking about at the time. [Filing] isn't reliable."

Users also may not file messages because they are concerned about failure to remember where information has been filed. Failure can have severe consequences especially if the message requires action:

"I don't know where to put it. And .. by making a wrong decision, I could really forget about it..."

¹This measures the number of messages that are part of extended conversations, rather than the total number of conversational threads. It underestimates the proliferation of threaded messages, because it does not detect the message that originated the thread. It also occasionally overestimates: people report taking the last message by another person and simply replying, to avoid having to regenerate another user's mailing address, and on rare occasions forgetting to modify the subject line when they do this.

²An exception occurs when users create a folder for an ongoing conversational thread, for which there is a large number of daily incoming messages. There seem to be two reasons for this: the first is the pressing need to remove the multiple interactions "cluttering up" the inbox, and the second is that the frequency of these messages means that the user will be *reminded* of the fact that the conversation is still in progress.

Another reason for not filing is that users want to postpone their judgement in order to determine the value of information. Users do not want to create archives containing information that later turn out to be useless or irrelevant. The strategy here is to wait and see the extent to which subsequent events indicate a message is valuable.

"I'm reluctant to archive junk. ... I know that the consequence of archiving junk is to make it that much harder to find the good stuff ... in the archive...Especially if information seems like it eventually will be overcome by events, I'd be very loath to move it into a [folder]. I'd be more likely to kind of hold it in my "in" box"

Folders may also not be useful after they are constructed. One problem is that users may not be able to remember folder labels, especially after a time has elapsed: *"if it's sort of older stuff, the category names are not going to mean anything to me any more"*. Users experienced special problems when they had large numbers of folders. They had to remember the definition of each when filing and to be careful not to introduce duplication by creating new folders that were synonymous with pre-existing ones. Duplication detracts from their use in retrieval.

In addition, folders can be *too small* to be useful. A major aim of filing is to reduce the huge number of undifferentiated inbox items into a relatively small set of folders each containing multiple related messages. Filing is clearly not successful if the number of messages in a given folder is small: if a folder contains only one or two items, then its existence has not significantly reduced the complexity of the inbox, nor gathered together significant amounts of related material. However our data show that filing often fails: on average 35% of users' folders contain only one or two items. Furthermore, not only do these tiny "failed folders" not significantly reduce the complexity of the inbox, the user has the dual overheads of (a) creating them in the first place, and (b) remembering multiple definitions every time there is a decision about filing an new inbox item.

The quantitative data reflect the problem of trying to remember multiple folder definitions. The larger the number of folders a user has, the more likely that person is to generate "failed folders" containing only one or two items ($r_{(16)} = 0.75, p < 0.001$). User statements bear this out:

"I wish I viewed creating a category as a lightweight activity. And for some reason I don't.... it seems like, you know the more of them I create, the harder it is to find any of them that are there!"

Folders can also fail because they are *too big*. When there are too many messages in a folder, it becomes unwieldy. It is difficult to find the relevant message in a large folder, the relationships between different messages in the folder become tenuous, so that one benefit of keeping them together is much reduced.

"So what happened was the size of the chunks associated with the categories got large. So now one key stroke would get me to a hundred things. So I really was no better off (filing information)"

To conclude, we have seen that users experience difficulties in creating folders. In addition, the returns for this effort may not be great: folders can be too large, too small or they may be too numerous for people to remember their individual definitions. As a consequence, folders may be of little use either for retrieval or for viewing related messages together. There is also a third problem:

filming information means that it is less available to *remind* users about that topic. Some users therefore finesse this problem: instead of filing incoming information, they simply leave it all in their inbox and use full-text search to find individual messages. We now examine users' strategies with respect to the problems of organising the inbox, so that it can be an effective method for managing ongoing tasks and conversations.

STRATEGIES FOR HANDLING EMAIL OVERLOAD

Given the dual problems of managing ongoing conversations and tasks, combined with the issue of filing, we identify three different user strategies, based on two criteria: (a) whether or not users currently use folders; (b) whether they "clean up" their inbox on a daily basis. This yields 3 strategies: *no filers* (no use of folders); *frequent filers* (folder users who try and clean up their inbox daily); *spring cleaners* (folder users who clean up their inbox only periodically). The data associated with each are shown in Table 1.

No filers: made no current use of folders (mean 11.33), but relied on full-text search to find information. Their folders were historic remnants from when two of the *no-filers* still filed. As a consequence of not filing, their inboxes were huge (3093.5 items, making up 95% of all their email). Their inboxes were overloaded: they included a large numbers of conversational threads (mean, 288). More significantly, over half of their inbox was old information that arrived more than 3 months ago. Their strategy for reducing the size of the overloaded inbox was periodic purges in which they deleted large numbers of old items or copied them to a separate independent archive. Four of the six *no filers* were managers.

Frequent filers: made strenuous attempts to minimise the numbers of inbox messages. They made daily passes through their inbox filing or deleting its contents. Their inboxes were relatively small, containing only 43.4 items, which was a very small percentage (5%) of the total number of mailbox messages. In addition, the inbox consisted almost exclusively of new items (90% were less than a month old, and only 5% were older than 3 months), and it was almost devoid of conversational threads (mean 3.6). They made frequent use of folders, and were relatively successful in their use of these, with only 21% being "failed folders". The five frequent filers included both the administrative assistants, but only one manager.

Spring Cleaners: dealt with the overloaded nature of their inboxes by intermittent *clean-ups* - normally every 1-3 months. They made extensive use of folders, even though this was often unsuccessful, as evidenced by the fact that over half of their folders "failed". They also had large overloaded inboxes (mean, 1492.3), containing large numbers of conversational threads (mean, 258). Over 40% of their inbox messages were more than 3 months old. Four of the seven *spring cleaners* were managers.

Impact, Overhead And Choice Of Strategies

All quantitative results are summarised in Table 1. In the remainder of the section, this data will be used to explore the impact of strategy choice. We discuss the trade-offs between strategies, and look at why users choose different strategies by examining the relationship between (a) strategy choice; and (b) factors such as job function and incoming message volume.

Strategy	N	# Inbox items	Total # items	Inbox as % of total mailbox	Old inbox items (%inbox> 3month old)	New inbox items (%inbox< 1month old)	# Inbox conversation threads	# Folders	Failed folders (# folders with < 3 items)	Daily # messages received
No filers	6	3093.5	3271.1	95.25	51.58	11.78	287.5	11.33	4.5	58
Spring Cleaners	7	1492.29	2818.71	51.02	40.15	24.22	257.86	61.43	30.85	45.71
Frequent filers	5	43.4	1062.2	4.96	4.75	90.34	3.6	70.6	16.6	42

Table 1 - The three strategies for managing email

The three strategies differ in terms of their benefits and maintenance costs. A major advantage of the *frequent filing* strategy is that the inbox can function as a genuine "to do" list containing a small number of outstanding messages relating to key current tasks, rather than having these items interspersed with huge numbers of conversational threads or unfiled old messages. It is also noteworthy that the average number of items in a frequent filer's inbox (43) can fit in two screens, and frequent filers report that it is important for them to be able to see items in this way. They are able to capitalise on the fact that when they view incoming unread email, they should be reminded of the majority of their outstanding tasks, because these should be visible or immediately accessible in the inbox.

"I don't have any other system, that keeps track of an e-mail message that needs a response usually the next day, hopefully it's still sort of near the bottom of the [inbox], I will see it when I look at new mail messages, so it won't get scrolled off the screen".

"But I live in the inbox. And that is kind of my to-do list. I'll keep things in there ... there's probably about twenty or thirty in there now of things that I want to keep like in my frontal lobes, that I have to deal with"

Frequent filers are effective in their use of folders, experiencing fewer "failed folders". This may be because of frequency with which they file enables them to remember the label definition and contents of each folder. However, despite the benefits of the clean "to do" list, opportunistic reminding, and the availability of current projects, there are major costs to this strategy. It requires significant maintenance: users have to make frequent passes through the inbox, filing and removing discharged items.

".. after I read the new day's mail... I go back to the whole "in" box, right back. And there's almost like a sifting that keeps happening, where the less pressing ones start ageing. And it gets to the point where I say, "I'm either not going to do something about this" And I just delete them".

It may be that frequent filing is only possible for lower volumes of incoming email, and for job specifications which do not require users to be away from their desk for long periods of time. Workers such as managers receiving higher volumes with less time to process email may not be in a position to exploit frequent filing.

The *no filing* strategy stands in direct contrast to *frequent filing*. Here users make few attempts to reduce the complexity of their inbox. When possible they answer messages as they receive them, but they

seldom review the inbox for outstanding undischarged messages. The fact that their inbox is cluttered with threads, as well as partially read and unread messages, means that outstanding tasks are not easily visible and are quickly displaced and scroll out of sight. Opportunistic reminding and task tracking are therefore unlikely to occur³. Users of this strategy admit that the clutter in their inboxes results in important tasks sometimes being overlooked.

"When you're dodging all this other stuff it's hard to pull out what could potentially be pretty important... like anyone else who may have that volume.... after a couple of days you're not going five screens up anymore. You're just looking at your current screen or maybe one more. Who knows what you've missed. I was on vacation for two weeks. Who knows what ... passed through ... e-mail during that time.... I saw a lot of it but I ... let a lot go by. I find that the most aggravating thing is when the inbox starts to grow and I don't know what to do with it necessarily".

When high volume of incoming email is accompanied by large amounts of time away from their desks, this may reduce the likelihood of *no-filers* constructing elaborate filing systems or engaging in extensive periodic clean-ups. The following *no-filer* described why he had abandoned any attempt to manage his inbox.

"Because what I used to do was use [spring cleaning] as a way as organizing and reviewing and catching anything that was falling off the end of the earth. I've given up on it. ... where am I going to get that time? If I wake up at three AM, and I've got nothing else to do, that's when I'm going to do it".

The *spring cleaners* are intermediate to the two other strategies. As with *no-filers*, as their inbox gets large, its size and complexity makes it ineffective as a "to do" list. The fact that it is usually cluttered with threads and unread messages means it is poor for task management, so opportunistic reminding is unlikely to occur. Furthermore the inbox was perceived to be of little archival use.

"it might as well be deleted as buried in this pile of junk E-mail may have value, but I will never avail myself in its current form, in this mail file. And so, it might as well be gone as sitting there, because either way I don't have it. It's not at my disposal and not usable".

Spring Cleaners have very strong feelings about the disorder of their inboxes: they use terms like "disgust" to describe their reactions to their inbox and are motivated by "seizures" to clean up. However the fact that they do occasionally go through the inbox means that

³One reason for the lack of task tracking may derive from *no filers*' organisational status: since they are usually managers, they tend to delegate tasks. They expect their employees to carry out the task, and report when it is done.

outstanding unprocessed messages are detected and can be replied to, even if these are sometimes late.

"So what I started to do was, either weekly, ... then monthly, I would go back to my mail, partly to categorize it and actually truthfully to catch things that I had just dropped the ball on"

This group seems to be less successful at creating useful folders than frequent filers. One possibility is that spring cleaners create folders infrequently, so that they forget folder definitions. Hence they may create duplicates of already existing folders. In terms of maintenance this strategy stands between the others: it does not require the daily efforts of *frequent filing*, but occasional clean-ups are required. This strategy choice may be explained by the inputs and workload of *spring cleaners*: they receive fewer messages and are less likely to be managers than *non-filers*, giving them more time to devote to managing their email.

We then tested these observations statistically. Because of our small subject pool, we were forced to combine data from both *spring cleaners* and *no filers*, in order to make comparisons. The analysis shows that *frequent filers* differ from the other strategies in a number of respects: they have smaller mailboxes ($t_{(16)} = 2.35, p < 0.05$), and smaller inboxes ($t_{(16)} = 3.94, p < 0.005$). Frequent filer's inboxes contain fewer inbox threads ($t_{(16)} = 3.99, p < 0.005$), and also tend to consist of newer items ($t_{(16)} = 2.41, p < 0.05$). Furthermore, there is a suggestion that they are more successful filers, with fewer "failed folders" ($t_{(16)} = 2.06, p = 0.058$).

Finally we looked at the impact on strategy choice, of factors such as organisational role and incoming volume of messages. We found only partial statistical evidence for the effects of role and volume. Managers were more likely to receive greater volumes of email ($t_{(16)} = 3.06, p < 0.005$). We then looked at whether managers were less likely to be frequent filers, given their higher volume of received email and greater time spent in meetings. Although only one frequent filer was a manager, there was no strong evidence of a direct relationship between strategy and status (*chi squared* (1 df) = 2.49, $p > 0.05$).

REDESIGNING EMAIL TO FIT ITS FUNCTIONS

There are both design and theoretical implications to these results. Although email was originally designed for *asynchronous communication*, the application is actually being used for multiple functions. Email therefore needs to be redesigned to support *filing* and *task management* as well as *asynchronous communication*. Our analysis of different users' strategies shows that both *non-filers* and *spring cleaners* experience problems with both *filing* and *task management*. These problems lead to backlogs of unanswered messages and "lost" information in archives. Furthermore, the group who experienced fewest problems, namely *frequent filers*, may only be able to operate successfully because of strenuous efforts to trim their inboxes. It is therefore important that we address overload. As email volumes continue to increase, even those users who are currently *frequent filers* may end up spending much of their energy in reading and responding to messages, leaving them with little time to maintain their inboxes, and folders. As a result, they may be forced to behave more like *spring cleaners* or even *non-filers*.

We now discuss possible techniques to support the three functions. We have shown that the inbox is often used as a place for incomplete

tasks, unfiled information and ongoing conversations. In all these cases, users preserve *working information* in the inbox both to keep it *available* and as a *reminder* that further actions are required. We have also seen, however, that opportunistic reminding is compromised when the number of inbox messages is too large, because messages scroll off the screen and remain unseen. A key technical requirement is therefore to *reduce inbox clutter* to allow visual reminding, but without compromising the *availability* of working information. We now present technical solutions for each email function addressing different ways of presenting and viewing the inbox to support both availability and reminding for working information.

Although email was originally designed for *asynchronous communication*, the current system has limitations in supporting this function. The key requirements for asynchronous communication are: (a) *threading* to support context regeneration and the management of conversational history, and (b) the ability to track the *status* of a conversation. Users want to avoid: scrolling back through large numbers of heterogeneous inbox messages to find all previous elements of a conversational thread; lost context when someone omits message history; forgetting who has the next turn in the conversational sequence.

How can we address these asynchronous communication problems? One solution to the problem of communication management automatically marks email messages from the same conversation using a common thread ID, allowing the user to collect related messages together, and trace back through conversations. The user would subsequently be able to *view by thread*. Viewing by thread allows a user to select any message, use that message to access all messages from that conversation, and hence view any message in its conversational context. This functionality is equivalent to having a single message containing the forwarded history of an entire conversation. Unlike a single message, viewing by conversation is not beset by the navigational problem of trying to follow a conversation that is many layers deep, where information may be buried within a single message. Viewing by thread provides several additional benefits. It helps determine *conversational status*: by looking at the last message in a thread, the user should be able to see whether they "owe" or are "owed" a response. Furthermore, it should be possible to file an entire thread, but *leave a representative message from that thread in the inbox*. This serves the purpose of *reducing inbox clutter*, even when users choose to copy themselves on every response. As we have seen with *frequent filing*, a representative message in an uncluttered inbox can *remind* the user that a conversation is in progress. When the conversation is concluded, the entire thread can also be archived or deleted as the user wishes.

What about *filing*? Given users' uncertainty about the value of much incoming information, they often end up with large numbers of incoming informational messages in their inbox. These documents are in a "holding pattern", while users attempt to determine their relevance and importance. In addition, our data show that filing may not be crucial for retrieval, because it may ultimately be superceded by full-text search. Nevertheless, users may want to cluster and view semantically related messages together, for example while they learn about a new topic.

How might we support this temporary buffering of incoming information? Information retrieval techniques could be used to cluster semantically related documents automatically, and the presentation of these clustered documents might be analogous with *conversational threads*. Users might therefore reduce the clutter of their inboxes, by leaving one semantic category exemplar in their inbox as a reminder and filing the rest. As with *threads*, each incoming message could be viewed in the context of other (in this case *semantically*) related information. This may partly address the problem of *failed folders*, and it might also help users to decide the usefulness of an arriving message. Two provisos are necessary here. First incoming documents should not be "filed" before the user has seen them. User comments and their experience with email filters clearly indicated that "automatic filing" was not desirable: users wanted to be made aware of the arrival of incoming documents, otherwise they would be ignorant of their existence. Furthermore, users were concerned that automatic filing would mean they wouldn't know the folder in which a given message had been filed. Second, this semantic classification needs to be *dynamic*, given that the status of a document can be changed by the arrival of subsequent ones.

Finally, when we consider *task management*, it is clear that conversational *threading* and *semantic clustering* should reduce the amount of inbox clutter by having each conversation or folder represented by one inbox message. The consequent reduction of the number of visible items in the inbox should help users to more easily see their outstanding tasks, and hence support *reminding* and *tracking*. Keeping important things "in view" could also be helped by having the inbox temporally sequenced and having threads and folders gradually "decay" by scrolling off the screen if they have not shown recent activity. Two further requirements seem to be crucial for task management. The first is the ability to *mark* particular inbox items as requiring action. This marker should be highly visible, and it should be possible to view only the "action items" in the absence of threads or folders. Note that this is different from creating an action *folder*: the action items are not filed away, but remain visible in the inbox, to serve their reminding function. The second requirement is the ability to *program reminders*. A critical problem occurs with "action items" that either can't be done immediately or don't need to be done at once. Here it would be useful to program these items so they would *re-appear* as an action item, as the deadline approaches.

Turning to outstanding research issues, we need to test the generality of these results. Although NotesMail is representative of current technology, we must see how the findings extend to both different email systems and a broader user population. It would be of particular interest to investigate systems lacking some of the features in NotesMail, e.g. text search, and also systems that have sophisticated filtering mechanisms, which NotesMail does not possess. To what extent do different features affect handling strategies? Our results about the effects of email volume and job specification on email processing strategies also need to be replicated with larger numbers of subjects.

Lastly, this work supports the findings of recent theoretical studies pointing to the interactive nature of cognition, and the fact that

people use external environmental sources or artifacts to mediate cognition [7,9,10,11]. We saw that email folders function as an external archival memory store. More importantly, people frequently access the email inbox, which means that properly organised it can operate as a visual device for attention manipulation and reminding, as a way of extending working memory, and maintaining the context of ongoing activities in an interrupt-driven environment. The importance of visual reminding and need to manage working information are reinforced by recent research on the organisation of physical and electronic desktops [1,2,6]. Past work has focussed almost exclusively on long term storage [3]. Future theoretical work should address these new issues of attention and the management of ephemeral and working information.

ACKNOWLEDGEMENTS

Thanks to the users for participating, to Sal Mazzotta for writing the analysis macros, also to Marilyn Walker, John Patterson, Irene Grief, Sara Kiesler and the members of Egrpsem for comments.

REFERENCES

1. Barreau, D. and Nardi, B. Finding and retrieving information: file organization from the desktop. To appear in *SIGCHI Bulletin*.
2. Kidd, A. The marks are on the knowledge worker. In *Proceedings of CHI'94 Human Factors in Computing Systems*, 186-191, ACM Press, New York, 1994.
3. Kraut, R., Fish, R., Root, B., & Chalfonte, B. Informal communication in organizations. In R. Baecker (Ed.), *Groupware and Computer Supported Co-operative Work*, 287-314, Morgan Kaufman, 1992.
4. Lansdale, M. The psychology of personal information management. *Applied Ergonomics*, 19, 55-66, 1988.
5. Maes, P. Agents that reduce work and information overload. In *Communications of the ACM*, 37, 31-40, 1994.
6. Malone, T. How do people organize their desktops? Implications for the design of Office Information Systems. *ACM Transactions on Office Information Systems*, 1, 99-112, 1983.
7. Norman, D. The psychology of everyday things. Basic Books: New York, 1988.
8. Sproull, L., and Kiesler, S. *Connections*. MIT Press, Cambridge, Mass., 1991.
9. Walker, M. Limited attention and discourse structure. *Computational Linguistics*, in press.
10. Whittaker, S., Frohlich, & Daly-Jones, O. Informal workplace communication: what is it like and how might we support it? In *Proceedings of CHI'94 Human Factors in Computing Systems*, 130-137, ACM Press, New York, 1994.
11. Whittaker, S., & Schwarz, H. Back to the future: pen and paper technology supports complex group co-ordination. In *Proceedings of CHI'95 Human Factors in Computing Systems*, 495-502, 1995.

*The observation that users do not want "automatic filing" has important implications for work advocating automatic filtering techniques [5].